PARALLEL UNSTRUCTURED VOLUME MESHING IN CFX-5

M. E. Braaten[a] and **M. J. Marchant**[b]

[a]GE Global Research
Niskayuna, NY 12309
braaten@crd.ge.com

[b]CFX, Harwell IBC
Didcot, OX11 0QR, UK
michael.marchant@cfx.aeat.com

A parallel advancing front algorithm for generating tetrahedral volume meshes has been developed and implemented in the commercial CFD code CFX5. The purpose of the parallel algorithm is to allow mesh generation to be performed on the same cluster of machines used to run the parallel flow solver, removing serial meshing as a bottleneck to running very large simulations. Although the primary objective is to allow larger meshes to be made, the parallel algorithm also has impressive benefits in reducing the time required to make a large mesh. An unstructured tetrahedral mesh with more 10 million nodes and 60 million elements has been successfully generated in under 2 hours on 27 processors of an SGI Origin 2000.

The paper outlines the basic procedures used in the existing CFX unstructured mesh generator to create hybrid tetrahedral / prismatic meshes. These include the interface to CAD, surface meshing, inflation of prism layers to resolve boundary layers and tetrahedral volume mesh generation. The parallel volume meshing algorithm is based on a domain decomposition of the advancing front algorithm pioneered by Peraire *et al* [1]. The basic serial algorithm already has a large number of enhancements that help insure high mesh quality, low memory usage and high element generation rates, it has seen commercial service in CFX5 for more than five years.

The parallel algorithm decomposes the geometry into a set of subdomains in which equal numbers of mesh points will be created. This decomposition is based on an estimate of the final mesh size generated from a novel procedure utilizing an Octree background mesh. The interiors of the subdomains are first meshed in parallel. This step represents the bulk of the computational effort and is almost fully parallelizable. The same well-proven parallel message passing libraries used in the CFX5 flow solver are used to perform the communication between processors. These libraries allow message passing via either PVM or MPI, and are suitable for shared memory MIMD machines as well as clusters of Unix workstations or Windows or Linux based PCs. Special care is taken at the end of this step to ensure a viable gap region between the subdomains, typically 2-4 cells wide. The gap region is currently meshed serially on the master processor, but this could be done recursively in parallel if the mesh and its corresponding gap region were large enough.

Great care was taken to ensure a fully scalable algorithm. A master/slave paradigm is used. Although 2D surface data may be replicated on each processor, no processor ever sees any global 3D arrays. The master processor handles all I/O operations, and the mesh is written out in pieces to a output single mesh file. The memory requirements for the master processor are no larger than those of the slave processors.

Results from the parallel algorithm have been obtained for a large number of complex industrial applications. Mesh quality is comparable to that achieved with the original serial algorithm. The final parallel volume meshing algorithm is now available in the commercial CFD code CFX5.6.

References

[1] J. Peraire, M. Vahdati, K. Morgan and O.C. Zienkiewicz, “Adaptive remeshing for compressible flow computations,” *J. Comp. Phys.*, 72, p. 449-466, 1987.